

VISIT...

LANZAROTE
Caliente.COM

The Age of Neuroelectronics

Adam Keiper

Every so often, when some new scientific paper is published or new experiment revealed, the press pronounces the creation of the first bionic man—part human, part machine. Science fiction, they say, has become scientific reality; the age of cyborgs is finally here.

Many of these stories are gross exaggerations. But something more is also afoot: There is legitimate scientific interest in the possibility of connecting brains and computers—from producing robotic limbs controlled directly by brain activity to altering memory and mood with implanted electrodes to the far-out prospect of becoming immortal by “uploading” our minds into machines. This area of inquiry has seen remarkable advances in recent years, many of them aimed at helping the severely disabled to replace lost functions. Yet public understanding of this research is shaped by sensationalistic and misleading coverage in the press; it is colored by decades of fantastical science fiction portrayals; and it is distorted by the utopian hopes of a small but vocal band of enthusiasts who desire to eliminate the boundaries between brains and machines as part of a larger “transhumanist” project. It is also an area of inquiry with a scientific past that reaches further back in history than we usually remember. To see the future of neuroelectronics, it makes sense to reconsider how the modern scientific understanding of the mind emerged.

The Body Electric

The brain has been clearly understood to be the seat of the mind for less than four centuries. A number of anatomists, philosophers, and physicians had, since the days of the ancient Greeks, concluded that the soul was resident in the head. Pride of place was often given to the ventricles, empty spaces in the brain that were thought to be home to our intelligent and immaterial spirits. Others, however, followed Aristotle in believing that the brain was just an organ for cooling the body. The clues that suggested its true function—like the brain’s proximity to most of the sensory organs, and the great safety of its bony encasement—were noticed but explained away. This is an understandable mistake. After all, how could that custard-like unmoving mass possibly house something as sublime and

Adam Keiper is managing editor of The New Atlantis.

4 ~ THE NEW ATLANTIS

Copyright 2006. All rights reserved. See www.TheNewAtlantis.com for more information.

complex as the human mind? Likelier candidates were to be found in the heart or in the body's swirling, circulating humors.

The modern understanding of the brain as the mind's home originated with a number of seventeenth-century philosophers and scientists. Among the most important was the Englishman Thomas Willis, an early member of the Royal Society, an accomplished physician, and a keen medical observer. Willis and his colleagues carefully dissected countless human brains, gingerly scooped from the skulls of executed criminals and deceased patients. He described his anatomical findings in several books, most notably *The Anatomy of the Brain and Nerves*, which included lovely and meticulous drawings by Christopher Wren. Willis described in great detail the structure of the brain and the body's system of nerves. He assigned the nerves a critical new role in the control of the body, and considered their study worthy of a new word, *neurology*. Carl Zimmer—whose enjoyable book *Soul Made Flesh* tells the story of Willis and the age of intellectual ferment and social turmoil in which he lived—details Willis's understanding of the body's nerves:

When Willis and his friends looked at them through a microscope they saw solid cords with small pores, like sugar cane. Reaching back to his earliest days of alchemy, Willis found a new way to account for how this sort of nerve could make a body move. He envisioned a nervous juice flowing through the nerves and animal spirits riding it like ripples of light. The spirits did not move muscles by brute force but rather carried commands from the brain to the muscles, which responded with a miniscule explosion. Each explosion, Willis imagined, made a muscle inflate.

We know today that Willis was not far off the mark—although instead of a nervous juice, we now know that nerves transmit electrical signals. This was a discovery long in coming; even though electricity had been used in medicine off and on for millennia. In olden days, it was often obtained by rubbing stones like amber; the Romans made medical use of electric eels. Inventions in the eighteenth century made it much easier to store electric charges, and the medical use of electricity became commonplace. Ben Franklin treated patients with shocks. So did Jean-Paul Marat. So, too, did John Wesley, the Methodist; he eventually opened three clinics for electrical treatment in London. The rapid rise and broad acceptance of electrotherapy in the eighteenth century, as chronicled in Timothy Kneeland and Carol Warren's book *Pushbutton Psychiatry*, is astonishing; it was used in treating a wide range of mental and physical ailments.

Despite the exposure of several notorious quacks, electrotherapy only became more popular in the nineteenth century, reaching its zenith in the decades just before and after 1900. Today's lastingly controversial practice of electroshock therapy (also called electroconvulsive therapy, because of the seizures it induces) can be considered a latter-day descendant of the old electrotherapy.

The scientific study of electricity and the nervous system progressed in tandem with the electrotherapy craze. A few researchers early in the eighteenth century suggested that nerves might transport electricity produced in the brain, but this was all speculation until the 1770s when Luigi Galvani noticed the twitching that occurred when dead frog legs were touched by two different metals. By the 1840s, scientists had used sensitive instruments to measure the tiny currents of nerves and muscles, and in 1850, the great German physicist and physiologist Hermann von Helmholtz succeeded in measuring the speed at which electrical impulses traversed the nervous system. The impulses traveled much slower than anyone expected. This was because the electrical signals weren't transmitted with lightning speed like signals on a copper wire; instead, the impulses were propagated by a slower biochemical process discovered later.

Scientists were also coming to understand more fully the functional structure of the brain. The scrutiny of patients with sick or injured brains—like the famous case of Phineas Gage, the railroad foreman whose personality changed radically in 1848 after a spike accidentally blew through his head—suggested to anatomists that skills and behaviors could be linked to specific brain locations. These clinical discoveries were complemented by laboratory research. At the beginning of the nineteenth century, Galvani's nephew, Giovanni Aldini, showed that electrical shocks to the brains of dead animals—and later dead criminals and patients—could produce twitches in several parts of their bodies. Decades later, other researchers continued this work more systematically, electrically shocking the brains of live animals to figure out which body parts were controlled by which spots on the brain.

By the 1890s, scientists had also worked out the cellular structure of the nervous system, using a staining technique that made it easier to see the fine details of the brain, the spinal cord, and the nerves. The individual nerve cells, called neurons, branch out to make connections with a great many other neurons. There are tens of billions of neurons in the adult human brain, meaning that there are perhaps a hundred trillion synapses where neurons can transmit electrical signals to one another.

The twentieth century brought great advances in psychopharmacology, and also a bewildering assortment of imaging technologies—X-rays, CT, PET, SPECT, MEG, MRI, and fMRI—that have made it possible to observe the living brain. Just as aerial or satellite photos of your hometown can convey a richer sense of ground reality than the battered old atlas in the trunk of your car, today's imaging technologies give a breadth of real-time information unavailable to the neural cartographers of a century ago. The latest research on brain-machine interfaces relies upon these new imaging technologies, but at its core is the basic knowledge about the nervous system—its electrical signals, its localized nature, and its cellular structure—already discovered by the turn of the last century. The only thing missing was a way of getting useful information directly out of the brain.

Brain Waves and Background Noise

In the 1870s, Richard Caton, a British physiologist, began a series of experiments intended to measure the electrical output of the brains of living animals. He surgically exposed the brains of rabbits, dogs, and monkeys, and then used wires to connect their brains to an instrument that measured current. "The electrical currents of the gray matter appear to have a relation to its function," he wrote in 1875, noting that different actions—chewing, blinking, or just looking at food—were each accompanied by electrical activity. This was the first evidence that the brain's functions could be tapped into directly, without having to be expressed in sounds, gestures, or any of the other usual ways.

Several years passed before others replicated Caton's work (in some cases, without awareness of his precedence), but even then, almost no one took notice. There was no easy way to keep records of the constant changes in their measurements of animal brain activity, so these early experimenters had to draw pictures of the activity their instruments measured. Only by 1913 did anyone manage to make the first crude photographic records of brain electrical measurements.

It wasn't until the 1920s that a researcher—German psychiatrist Hans Berger—first measured and recorded the electrical activity of human brains. As a young man, Berger had experienced an odd coincidence that led him to believe in telepathy. This influenced his decision to study the connection between mind and matter, and led him to research "psychic energy." He spent decades trying to measure the few quantifiable brain processes involving energy—the flow of blood, the transfer of heat,

and electrical activity—and attempting to link those physical processes to mental work. Electrical measurement was of special interest to Berger, and whenever he could get away from his family, his patients, and his many administrative obligations, he would sequester himself in a laboratory from which he barred colleagues and visitors.

The great difficulty facing Berger was to isolate the brain's activity amidst the electrical cacophony of the body and through the thick obstruction of the skull using instruments that were barely sensitive enough for the task. His first successful measurements were on patients with fractures or other skull injuries that left spots with less bone in the way. (The recently concluded war had something to do with the availability of such patients.) Slowly improving his instrumentation through years of frustrating trial and error, by 1929 Berger was finally reliably producing records of the brain activity of subjects with intact skulls, including his son and himself. He coined the word *electroencephalogram* for his technique, and published more than a dozen papers on the subject.

Berger's electroencephalograms (EEGs) represented the brain's electrical activity as complicated lines on a graph, and he tried to discriminate between the various underlying patterns that made up the whole. He believed that certain recurring wave patterns with discernible shapes—which he called alpha waves, beta waves, and so forth—could be linked to specific mental states or activities. A few years passed before other researchers took notice of Berger's work; when they finally did, in the mid-1930s, there was rapid progress in picking apart the patterns of the EEG.

One early breakthrough was the use of the EEG to locate lesions on the brain. Another was the discovery of a particular wave pattern—an unmistakable repeating “spike-and-dome”—connected to epilepsy. This pattern was so pronounced that the United States Army Air Corps began using EEGs during World War II to screen out pilots who might have seizures. There was even some discussion about the possible use of EEG as a eugenic tool—akin to the way genetic counseling is sometimes used today. “Couples who believe in eugenics may yet exchange brain-wave records and consult an authority on heredity before they marry,” said one 1941 *New York Times* article. “A man and a woman who may be outwardly free from epilepsy but whose brain waves are of the wrong shape and too fast are sure to have epileptic children.”

That term—“brain waves”—actually antedates the EEG by several decades. It was used as early as the 1860s to describe a “hypothetical telepathic vibration,” according to the *Oxford English Dictionary*. As the public

slowly came to learn about the wavy lines of the EEG, the term donned a more respectable scientific mantle, and newspaper articles during the 1940s sometimes spoke of the EEG as a “brain-wave writer.” Perhaps Berger, whose initial impetus for neurological research was his own interest in telepathy, would have been amused by the terminological transition from fancy to fact.

What that etymological shift somewhat obscures, though, is that EEG is most assuredly *not* mind-reading. The waves of the EEG do not actually represent thoughts; they represent a sort of jumbled total of many different activities of many different neurons. Beyond that, there remains a great deal of mystery to the EEG. As James Madison University professor Joseph H. Spear recently pointed out in the journal *Perspectives in Science*, there remains a “fundamental uncertainty” in EEG research: “No one is quite certain as to what the EEG actually measures.” This mystery can be depressing for EEG researchers. Spear quotes a 1993 EEG textbook that laments the “malaise” and “signs of pessimism, fatigue, and resignation” that electroencephalographers evince because of the slow theoretical progress in their field.

Here’s one way to think about the great challenge facing these researchers: Imagine that your next-door neighbor is having a big dinner party with some foreign friends who speak a language you don’t know. Your neighbor’s windows are closed, his curtains are drawn shut, and his stereo is blasting loud music. You aren’t invited, but you want to know what his guests are talking about. Well, by listening intently from outside a window to laughter and lulls and cadences, you can probably figure out whether the conversation is friendly or angry, whether the partygoers are bored or excited, maybe even whether they are talking about sports or food or the weather. You would try hard to ignore the sounds of the stereo. Maybe you would call your neighbor on the telephone and ask him to tell you what his foreign friends are talking about, although there is no guarantee that his description will be accurate. You might even try to affect the conversation—maybe by flashing a bright light at the window—just to see how the partygoers react.

EEG research is somewhat similar. Researchers try to look for patterns, provoke responses, and tune out background noise. Since the 1960s, they have been aided in their work by computers, which use increasingly sophisticated “signal processing” techniques to filter out the din of the party so that an occasional whisper can be heard. By exposing a patient to the same stimulus again and again, investigators can watch for repeating reactions. Using these methods, researchers have been able to go beyond

the old system of alpha, beta, and delta waves to pick out more subtle spikes, dips, and bumps on the EEG that can be linked to action, reaction, or expectation.

Since our brains produce these tiny signals without conscious command, it is not surprising that there is interest in exploiting some of these signals to “read the mind” in the same way that pulse, galvanic skin response, and other indicators are used in lie detectors. In fact, a neuroscientist named Lawrence A. Farwell has gotten a great deal of press in the last few years for marketing tests that rely heavily on an EEG wave called P300. The P300 wave, which has been researched for decades, has been called the “Aha!” wave because it occurs a fraction of a second after the brain is exposed to an unexpected event—but before a conscious response can be formulated and expressed. Farwell, who spent two years working with the CIA on forensic psychophysiology, has founded a company called Brain Fingerprinting Laboratories. The company’s website boasts of helping to free an innocent man serving a life sentence for murder in an Iowa prison: the P300 test “showed that the record stored in [the convict’s] brain did not match the crime scene and did match his alibi.” The district court admitted the P300 test as evidence. When a key accuser was “confronted with the brain-fingerprinting evidence,” he changed his story, and the convict was freed after two decades in prison. The use of the P300 test was not without controversy, however: Among the witnesses testifying against its admissibility in court was Emanuel E. Donchin, the preeminent P300 researcher, and a former teacher of and collaborator with Farwell. Donchin has repeatedly said that much more research and development needs to be done before the technique can be used and marketed responsibly.

The criticism hasn’t slowed Farwell, however. His company’s website describes how the P300 wave helped put a guilty man behind bars for a long-unsolved murder; it tells of Farwell’s fruitless eleventh-hour efforts to save a murderer from execution because the P300 test supposedly indicated his innocence; and it discusses how the P300 test might be used to diagnose Alzheimer’s disease and to “identify trained terrorists.” While there is little reason to believe the P300 test will be so used—after all, the traditional means of determining dementia or identifying terrorists seem simpler—it is conceivable that the P300 test or something similar will someday become more refined and more widely accepted, replacing older, and notoriously unreliable, lie-detection technology.

The P300 test relies on one of several electrical signals that the conscious mind generally cannot control. Yet one of the major applications

of EEG has been to exert more conscious control upon the unconscious body. “Biofeedback” is the name of a controversial set of treatments generally classified alongside acupuncture, chiropractic, meditation, and other “alternative therapies” that millions of people swear by even though the medical establishment frowns its disapproval. Biofeedback treatments that use EEG are sometimes called “neurofeedback,” and they generally work something like this: A patient wears electrodes that connect to an EEG and is given some kind of representation of the results in real-time. This is the feedback, which can be a tone, an image on a screen, a paper print-out, or something similar. The patient then tries to change the feedback (or maintain it, depending on the purpose of the therapy) by thinking a certain way: clearing his mind, or concentrating very hard, or imagining a particular activity.

This may sound absurd—and indeed, much of the literature about neurofeedback is quite kooky, rife as it is with mystical mumbo-jumbo—but evidently enough people are interested to sustain a small neurofeedback industry. Steven Johnson hilariously described visits to several neurofeedback companies in his 2004 book *Mind Wide Open*. First, he meets with representatives from The Attention Builders, a company whose Attention Trainer headset and software is intended for children with attention-deficit disorder. The company has “concocted a series of video games that reward high-attention states and discourage more distracted ones,” Johnson writes. “Start zoning out while connected to the Attention Trainer software, and you’ll see it reflected on the screen within a split second.” Then he visits Braincare, a neurofeedback practice in New York, and uses its similar system to control an onscreen spaceship—“and once again I find that I can control the objects on the screen with ease.” Next Johnson visits a California-based practice run by the Othmers, a couple who first encountered neurofeedback in 1985 when they were looking for some way to help their neurologically-impaired son control his behavior. Soon, the entire Othmer family was using neurofeedback therapy—mother for her hypoglycemia, brother for his hyperactivity, and father for a head injury. So convinced were the Othmers of the efficacy of neurofeedback that they made a career of it. Johnson describes his experience with the Othmers’ system for training patients to control their mental “mode”:

Othmer suggests that we start with a more active, alert state. She hits a few buttons, and the session begins. I stare at the Pac-Man and wait a few seconds. Nothing happens. I try altering my mental state, but

mostly I feel as though I'm altering my facial expression to convey a sense of active alertness, as though I'm sitting in the front row of a college lecture preening for the professor. After a few seconds, the Pac-Man moves a few inches forward, and the machine emits a couple of beeps. I don't really feel any different, but I remember Othmer's mantra—"be pleased that it's beeping"—and so I try to shut down the part of my brain that's focused on its own activity, and sure enough the beeping starts up again. The Pac-Man embarks on an extended stroll through the maze. I am pleased.

Johnson's experience, like similar anecdotes from neurofeedback patients, demonstrates just how difficult it can be, especially for novices, to control the sorts of brain activity that an EEG picks up. And even though biofeedback therapy is unlikely to migrate from the fringes to the mainstream of medical acceptability, we shall see that essentially the same EEG technique is now being pursued by many researchers attempting to build brain-machine interfaces. As one of the leading brain-machine interface researchers told the *New Yorker* in 2003, his work could rightly be called "biofeedback"—but he doesn't want anyone to confuse it with that "white-robed meditation crap."

Into the Brain, Into the Mind

While EEG provides a kind of confused, collective sense of the brain's electrical activity, there is a much more direct way to tap into the brain: stick an electrode into it. This approach allows not only for the measurement of electrical activity in parts of the brain, but also for the direct electrical stimulation of the brain.

The forerunners of today's brain implants can be found in the nineteenth century efforts to map different brain functions by shocking different parts of the brains of anesthetized or restrained animals. These efforts continued for decades, yielding a picture of the brain that was both increasingly detailed and stupefyingly complex. But this great body of work revealed very little about the brain's electrical activity during normal behavior, since there were practically no attempts to put electrodes in the brains of animals that weren't drugged or restrained. Stanley Finger's *Origins of Neuroscience* tells of a little-known German professor named Julius R. Ewald "who put platinum 'button' electrodes on the cortex of [a] dog in 1896," then walked the dog on a leash and "stimulated its brain by connecting the wires to a battery." Finger notes that Ewald "did not write up his work in any detail, but a young American who visited

Germany extended Ewald's work and then published a more complete report of these experiments" in 1900.

The first scientist to use brain implants in unrestrained animals for serious research was a Swiss ophthalmologist-turned-physiologist named Walter Rudolf Hess. Starting in the 1920s, Hess implanted very fine wires into the brains of anesthetized cats. After the cats awoke, he sent small currents down the wires.

This experiment was part of Hess's research into the autonomic nervous system, work for which he was awarded a Nobel Prize in 1949 (sharing the prize with the Portuguese neurologist Egas Moniz, the father of the lobotomy). In his Nobel lecture, Hess described how his stimulation of the animals' brains affected not merely their motions and movements, but also their moods:

On stimulation within a circumscribed area...there regularly occurs namely a manifest change in mood. Even a formerly good-natured cat turns bad-tempered; it starts to spit and, when approached, launches a well-aimed attack. As the pupils simultaneously dilate widely and the hair bristles, a picture develops such as is shown by the cat if a dog attacks it while it cannot escape. The dilation of the pupils and the bristling hairs are easily comprehensible as a sympathetic effect; but the same cannot be made to hold good for the alteration in psychological behavior.

In the decades that followed, a great many researchers began to use implanted brain electrodes to tinker with animal and human behavior. Three individuals are of particular interest: James Olds, Robert Heath, and José Delgado.

James Olds was a Harvard-trained American neurologist working in Canada when, in 1953, he discovered quite by accident that a rat seemed to enjoy receiving electric shocks in a particular spot in its brain, the septum. He began to investigate, and discovered that the rat "could be directed to almost any spot in the box at the will of the experimenter" just by sending a zap into its implant every time it took a step in the desired direction. He then found that the rat would rather get shocked in its septum than eat—even when it was very hungry. Eventually, Olds put another rat with a similar implant in a Skinner box wherein the animal could stimulate itself by pushing a lever connected to the electrode in its head; it pressed the lever again and again until exhaustion.

Thus was the brain's "pleasure center" discovered—or, as Olds came to describe it later because of its winding path through the brain, the "river of reward." It was soon established that other animals, including humans, have

similar pleasure centers. Countless researchers have studied this area over the years, but perhaps none more notably than Robert Galbraith Heath. A controversial neuroscientist from Tulane University in New Orleans, Heath in the early 1950s became the first researcher to actually put electrodes deep into living human brains. Many of his patients were physically ill, suffering from seizures or terrible pain. Others came to him by way of Louisiana's state mental hospitals. Heath tried to treat them by stimulating their pleasure centers. He often met with remarkable success, changing moods and personalities. With the flip of a switch, murderous anger could become light-heartedness, suicidal depression could become contentment. Conversely, stimulating the "aversive center" of a subject's brain could induce rage.

By the 1960s, Heath had begun experimenting with self-stimulation in humans; his patients were allowed to trigger their own implants in much the same way as Olds's rats. One patient felt driven to stimulate his implant so often—1,500 times—that he "was experiencing an almost overwhelming euphoria and elation, and had to be disconnected, despite his vigorous protests," Heath wrote. The strange story of what happened to that patient next, in an experiment so thoroughly politically incorrect that it would never be permitted today, is recounted in Judith Hooper and Dick Teresi's outstanding book *The Three-Pound Universe*:

[The patient] happened to be a schizophrenic homosexual who wanted to change his sexual preference. As an experiment, Heath gave the man stag films to watch while he pushed his pleasure-center hotline, and the result was a new interest in female companionship. After clearing things with the state attorney general, the enterprising Tulane doctors went out and hired a "lady of the evening," as Heath delicately puts it, for their ardent patient.

"We paid her fifty dollars," Heath recalls. "I told her it might be a little weird, but the room would be completely blacked out with curtains. In the next room we had the instruments for recording his brain waves, and he had enough lead wire running into the electrodes in his brain so he could move around freely. We stimulated him a few times, the young lady was cooperative, and it was a very successful experience." This conversion was only temporary, however.

Another brain-implantation pioneer, José Manuel Rodríguez Delgado, described how he induced the same effect in reverse: when a particular point on a heterosexual man's brain was stimulated, the subject expressed doubt about his sexual identity, even suggesting he wanted to marry his male interviewer and saying, "I'd like to be a girl."

That experiment is described in Delgado's riveting 1969 book, *Physical Control of the Mind*. A flamboyant Spanish-born Yale neuroscientist, Delgado, like Heath, began exploring in the 1950s the electrical stimulation of the reward and aversion centers in humans and animals—what he called “heaven and hell within the brain.” Like Heath, Delgado tells stories of patients whose moods shifted after their brains were stimulated—some becoming friendlier or flirtatious, others becoming fearful or angry. He describes artificially inducing anxiety in one woman so that she kept looking behind her and said “she felt a threat and thought that something horrible was going to happen.” In other patients, Delgado triggered hallucinations and déjà vu.

Delgado invented a device he called the “stimoceiver,” an implant that could be activated remotely by radio signal. The stimoceiver featured prominently in the experiment for which Delgado is best known, in which he played matador, goading a bull into charging him, only to turn off the bull's rage with a click of the remote control at the last instant. The bull had of course had a stimoceiver implanted in advance.

This is just one of a great many bizarre animal experiments detailed in Delgado's brilliant, absurd, coldhearted, sickening book. A weird menagerie of animals with brain implants is shown in the book's photographs. One little monkey is electrically stimulated so that one of its pupils dilates madly. Friendly cats are electrically provoked to fight one another. Chimpanzees “Paddy” and “Carlos” have massive implants weighing down their heads. One rhesus monkey is triggered 20,000 times so that the scientists can observe a short ritual dance it does each time; another loses its maternal instinct and ignores its offspring when triggered; yet another is controlled by cagemates that have learned that pressing a lever can bring on docility.

The creepiest revelation of Delgado's book is how easily the brain can be fooled into believing that it is the source of the movements and feelings actually induced by electrical implants. For instance, a cat was stimulated in such a way that it raised one of its rear legs high into the air. “The electrical stimulation did not produce any emotional disturbance,” Delgado writes, but when the researchers tried to hold down the cat's leg, it reacted badly, “suggesting that the stimulation produced not a blind motor movement but also a desire to move.” A similar effect was noticed with humans, as in the case of a patient who was stimulated in such a way that he slowly turned his head from side to side:

The interesting fact was that the patient considered the evoked activity spontaneous and always offered a reasonable explanation for it. When asked, “What are you doing?” the answers were, “I am looking for my

slippers,” “I heard a noise,” “I am restless,” and “I was looking under the bed.” In this case it was difficult to ascertain whether the stimulation had evoked a movement which the patient tried to justify, or if an hallucination had been elicited which subsequently induced the patient to move and to explore the surroundings.

When other patients had their moods suddenly shifted by stimulation, they felt as though the changes were “natural manifestations of their own personality and not...artificial results of the tests.” One might reasonably wonder what such electrical trickery and mental manipulability suggest about such concepts as free will and consciousness. To Delgado, these are but illusions: he speaks of the “so-called will” and the “mythical ‘I.’” It is not surprising, given his totally physicalist views, that Delgado should end his book with a call for a great program of researching and altering the human brain with the aim of eliminating irrational violence and creating a “psychocivilized society.”

There was admittedly some interest in such ideas for a short while in the late 1960s and early 1970s, primarily among those hoping to study and rehabilitate prison inmates by means of electrical implants. The chief byproduct of these efforts, it would seem, was the creation of a lasting paranoia about U.S. government plans to control the population with brain implants. (Do an Internet search for “CIA mind control” to see what I mean.) But in reality, of course, the vast social program Delgado envisioned never came to pass, and in some ways, research into the manipulation of behavior through electrical stimulation of the brain has not gone very far beyond where Delgado, Heath, and their contemporaries left it. Consider, for example, the brain-controlling implant that received the most attention in the past few years: a 2002 announcement by researchers at the SUNY Downstate Medical Center that they could control the direction that rats walk (through mazes, across fields, and so forth) by remotely stimulating the pleasure centers in their brains. The scientists claimed that this disturbing research might eventually have practical applications—like the use of trained rats in search-and-rescue operations. But the media excitement about these “robo-rats” obscured the fact that this remote-controlled rodent perambulation was barely an advancement over the work James Olds first did with rats a half-century ago.

The Brain Pacemaker

In his 1971 novel *The Terminal Man*, Michael Crichton imagined the first-ever operation to insert a permanent electrical implant into the brain of a

man suffering from psychomotor epilepsy. In the story, the patient's seizures and violent behavior are repressed by jolts from the implant. Relying on the best available prognostications about how such futuristic technology could work, Crichton meticulously described every detail: the surgery to insert the forty-electrode implant; the implant's long-lasting power pack; the tiny computer inserted into the patient's neck to trigger the implant when a seizure was imminent; and the testing, calibration, and use of the implant. In true Crichton style, things go awry soon after the surgery and the patient runs away from the hospital and starts killing people.

Similar surgeries were being carried out in real life in the United States just a few years later. Perhaps the first was an operation to insert an implant designed by Robert Heath and his colleagues, a permanent version of the implants Heath had used in the previous decade. The patient was a mentally-retarded young man prone to fits of terrible violence. Some of the things Crichton predicted hadn't yet been developed—so, for example, Heath's real-life implant didn't have a tiny computer telling it when to zap the brain, it just zapped on a regular schedule, much as an artificial pacemaker sends regular electrical impulses into the heart. And instead of a small power pack under the skin, Heath's implant was connected by wire to a battery outside the skin.

The operation had the desired effect and the patient became sufficiently calm to go back home—until, as in Crichton's story, something went wrong and the patient abruptly became violent and tried to kill his parents. But unlike Crichton's implantee, who met with a bloody end, the real-life patient was captured and safely returned to the hospital, where Heath promptly discovered that the problem was caused by a break in the wires connecting the implant to its battery. The battery was reconnected and the patient went back home.

In all, Heath and his colleagues inserted more than seventy similar implants in the 1970s, with some of the patients seeing dramatic improvements and about half seeing "substantial rehabilitation," according to Hooper and Teresi.

In the years that followed, the study of brain implants stalled. Then, in the 1980s, after French doctors discovered that an electrode on the thalamus could halt the tremors in a patient with Parkinson's disease, researchers began to focus on the feasibility of treating movement disorders with electrical implants. Previously, the only treatments available to Parkinson's patients were drugs (like levodopa, which has a number of unpleasant side effects) and ablative surgery (usually involving either the intentional scarring or destruction of parts of the brain). Other hoped-

for cures, like attempts to graft dopamine-producing cells from kidneys or even fetuses into the brains of Parkinson's patients, weren't panning out. But subsequent research confirmed the French doctors' discovery, and a company called Medtronic—one of the first producers of cardiac pacemakers in the late 1950s—began work on an electrical implant for treating Parkinson's patients. After several years of clinical investigation, the Medtronic implant was approved in 1997 by the U.S. Food and Drug Administration (FDA) for use in treating Parkinson's disease and essential tremor; in 2003, it was approved for use in treating another debilitating movement disorder called dystonia.

The implant is often called a "brain pacemaker" or a "neurostimulator"; Medtronic uses brand names like Soletra and Kinetra. The treatment goes by the straightforward name "Deep Brain Stimulation." In the procedure, a tiny electrode with four contacts is permanently placed deep in the brain. It is connected by subcutaneous wire to a device implanted under the skin of the chest; this device delivers electrical pulses up the wire to the electrode in the brain. (Many patients are given two electrodes and two pulse generating devices, one for each side of the body.) The device in the chest can be programmed by remote control, so the patient's doctor can pick which of the four contacts get triggered and can control the width, frequency, and voltage of the electrical pulse. Patients themselves aren't given the same level of control, but they can use a handheld magnet or a remote control to start or stop the pulses.

The first thing that must be said about Deep Brain Stimulation is that it really works. There is always a risk of complications in brain surgery, and even when the operation is successful there is no guarantee that the brain pacemaker will bring the patient any relief. But more than 30,000 patients around the world with movement disorders have had brain pacemakers implanted, and a majority of them have apparently had favorable results. Some of the transformations seem miraculous. The *Atlanta Journal-Constitution* tells of Peter Cohen, a former lawyer whose dystonia robbed him of his livelihood and left him stooped, shaking, and often stretched on the floor of his home; less than two years after his operation, he was off all medication, walking about normally without attracting the stares of passersby, and hoping to resume his law career. A single mother told the *Daily Telegraph* of London how her tremors from Parkinson's abated after she received her implant; before the surgery she had been taking "forty or fifty pills a day," but afterwards she was off all medication and feeling "like a normal mum again." Tim Simpson, a top professional golfer, had to quit playing after a series of health problems, including essential tremor; after

he had his brain pacemaker implanted, his hand steadied and he has since returned to pro golf, according to a profile in the *Chicago Sun-Times*. The *San Francisco Chronicle* describes a family in which three generations have had the implants: a mother and her elderly father with essential tremor have gotten over their trembling, while her teenage son with dystonia has regained the ability to walk. Thousands of similarly treated patients have come to regain normal lives.

The second thing that must be said about Deep Brain Stimulation is that nobody knows how it works. There are many competing theories. Perhaps it inhibits troublesome neural activity. Or maybe it excites or regulates neurons that weren't firing correctly. Some researchers think it works at the level of just a few neurons, while others think that it affects entire systems of neurons. That it works is undeniable; how it works is a puzzle far from being solved.

Given the mysteriousness of Deep Brain Stimulation, it should come as no surprise that the implants seem to be capable of much more than just stopping tremors. According to various sources, scientists are investigating the use of the implants for treating epilepsy, cluster headaches, Tourette's syndrome, and minimally-conscious state following severe brain injury. What's more, since at least the 1990s it has been clear that the implants can affect the mind, and in the past few years they have been used experimentally to treat a few dozen cases of severe depression and obsessive-compulsive disorder—cases where several other therapies had failed. Establishing experimentally whether such treatments will work is tricky business, since there can be no animal tests for these mental illnesses and since it's all but impossible to conduct blind studies for brain implantation surgeries. But the evidence from the small pool of such patients treated so far seems to show that several have been helped, although none has been cured.

The evidence also suggests that these implants affect mood and mind more subtly than those used by Delgado and Heath more than a generation ago. Consider this 2004 testimony from G. Rees Cosgrove, a Harvard neurosurgeon, to the President's Council on Bioethics:

So we have four contacts in [the brain], and Paul Cosyns, who is one of the investigators in Belgium, relates this very wonderful anecdote that one of the patients [he] successfully treated has, you know, their four contacts, and she says, "Well, Dr. Cosyns, when I'm at home doing my regular things, I'd prefer to have contact two [activated], but if I'm going out for a party where I have to be on and, you know, I'm going

to do a lot of socializing, I'd prefer contact four because it makes me revved up and more articulate and more creative."...

We have our own patient who is a graphic designer, a very intelligent woman on whom we performed the surgery for severe Tourette's disorder and blindness resulting from head tics that cause retinal detachments, and we did this in order to try and save her vision. The interesting observation was that clearly with actually one contact we could make her more creative. Her employer saw just an improvement in color and layout in her graphic design at one specific contact, when we were stimulating a specific contact.

These stories suggest that brain implants could be used intentionally to improve the mental performance of healthy minds with less imprecision than mind-altering drugs and less permanency than genetic enhancement. But that possibility is remote. For the foreseeable future, there is no reason to believe that any patient or doctor will attempt to use Deep Brain Stimulation with the specific aim of augmenting human creativity. The risks are too high and the procedure is too expensive. But even if the surgery were much safer and cheaper, we know so little about how these implants affect the mind that any such attempt would be as likely to dull creativity as to sharpen it.

More significant is the possibility that implants will, in time, move into the mainstream of treatment for mental illness. Not counting the thousands of motion-disorder patients with brain pacemakers, Deep Brain Stimulation has so far only been tried on a few severe cases of mental illness. But there is another technique that involves the stimulation of the vagus nerve in the neck; it is mainly used in the treatment of epilepsy, but for the past few years has been used in Canada and Europe to treat the severely depressed. In July 2005, the FDA approved it for use as a last-resort treatment for depression. According to a recent article in *Mother Jones*, Cyberonics, the company that makes the vagus nerve stimulator, "has hired hundreds of salespeople to chase after the 4 million treatment-resistant depressives that the company says represent a \$200 million market—\$1 billion by 2010." You may even have seen some of the Cyberonics direct-to-consumer advertisements online. Consider this recipe for a new industry: ambitious companies eager to break into a new market, vulnerable consumers looking for pushbutton relief, and growing ranks of neurosurgeons with implant experience. How long before patients pressure their doctors to prescribe an implant? How long before the defining-down of "last resort"? How long until brain stimula-

tion becomes the neuromedical equivalent of cosmetic surgery—drawing upon real medical expertise for non-medical purposes?

Of course, this may never come to pass. Implants may stay too dangerous for all but the worst cases—or implant therapy for mental illness might be outpaced and obviated by improved psychopharmacological therapies. But there are those who would like to see brain implants become a matter of choice, even for the healthy. David Pearce, a prominent British advocate of transhumanism, has argued that implants in the brain's pleasure centers should be one technological component of a larger project to abolish suffering. His musings on this subject are outlined in an intriguing, if laughably idealistic, manifesto called "The Hedonistic Imperative." (On one of his websites, Pearce offers this recent quote purportedly from the Dalai Lama: "If it was possible to become free of negative emotions by a riskless implementation of an electrode—without impairing intelligence and the critical mind—I would be the first patient.") Pearce's idealism may seem, on the surface, to be the antithesis of Delgado's dreams of an imposed "psychocivilized society." But they are of a piece. Enamored of the possibilities new technologies open up, unsatisfied with given human nature, and unburdened by an appreciation for the lessons of history, they both forsake reality for utopia.

One Letter at a Time

The most compelling research being done on brain-machine interfaces is as far from utopia as can be imagined. It is in the hellish reality of a trapped mind.

Modern medicine has made it possible to push back the borders of "the undiscover'd country" so that tiny premature babies, the frail elderly, and the gravely sick and wounded can live longer. One consequence has been the need for new categories that would have gone unnamed a century ago—"brain death" (coined 1968), "persistent vegetative state" (coined 1972), "minimally conscious state" (coined 2002), and so on. Perhaps the most terrifying of these categories is "locked-in syndrome" (coined 1966), in which a mentally alert mind is entrapped in an unresponsive body. Although the precise medical definition is somewhat stricter, in general usage the term is applied to a mute patient with total or near-total paralysis who remains *compos mentis*. The spirit is willing, but the flesh is weak. The term is sometimes used to describe a patient who retains or regains some slight ability to twitch and control a finger or limb, but locked-in patients can generally only communicate by blinking or by moving their

eyes—movements that must be interpreted either by a person or an eye-tracking device. Sometimes they lose even that ability. Swallowing, breathing, and other basic bodily functions often require assistance. In a word, it is the greatest state of dependency an awake and sound-minded human being can experience.

Locked-in syndrome can develop inexorably over time as the result of a degenerative disease like ALS (Lou Gehrig's disease), or it can be the sudden result of a stroke, aneurysm, or trauma. Misdiagnosis is a frequent problem; there have been documented cases of locked-in patients whose consciousness went unnoticed for years; in more than a few cases, locked-in patients have reported the horror of being unable to reply when people within earshot debated disconnecting life support.

Statistics are nonexistent, but there are surely thousands, and perhaps tens of thousands, of locked-in patients (depending on how broadly the term is defined). Their plight has received attention in recent years partly because of a number of books and articles written by locked-in patients, painstakingly spelling out one letter at a time with their eyes. A Cornell student paralyzed by a stroke at age 19 described her fears and frustrations in her 1996 book, *Locked In*. A former publishing executive in France defiantly titled his 1997 memoir of locked-in syndrome *Putain de silence* (*F***ing Silence*; the English version was given the sanitized title *Only the Eyes Say Yes*). A young rugby-playing New Zealander left locked-in by strokes described in a 2005 essay in the *British Medical Journal* how he “thought of suicide often” but “even if I wanted to do it now I couldn’t, it’s physically impossible.” By far the most famous account of a locked-in patient is *The Diving Bell and the Butterfly*, a bestseller written by French magazine editor Jean-Dominique Bauby. He spent less time locked-in than the other patient-authors—his stroke was in December 1995, he dictated his book in 1996, and he died two days after it was published in 1997—but his account is the most poignant and poetic. The book’s title refers to his body’s crushing immobility while his mind remains free to float about, flitting off to distant dreams and imaginings. He describes the love and the memories that sustain him. And he tells of the times when his condition seems most “monstrous, iniquitous, revolting, horrible,” as when he wishes he could hug his visiting young son.

Brain-machine interfaces are likely to make it easier for patients with locked-in syndrome to communicate their thoughts, express their wishes, and exert their volition. Experimental prototypes have already helped a few locked-in patients. With sufficient refinement, brain-machine interfaces may also make life easier for patients with less total paralysis—although

for years to come, any patient retaining command of a single finger will likely have more control over the world than any brain-machine interface can provide.

The concept behind this kind of brain-machine interface is simple. We know that electrical signals from brains can be detected by electrode implants or by EEG. But what if the signals were sent to a machine that does something *useful*? Although most of the serious research in this area goes back only to the 1980s, there are some earlier examples. Perhaps the first is a 1963 experiment conducted by the eccentric neuroscientist and roboticist William Grey Walter. Patients with electrodes in their motor cortices were given a remote control that let them advance a slide projector, one slide at a time. Grey Walter didn't tell the patients, though, that the remote control was fake. The projector was actually being advanced by the patients' own brain signals, picked up by the electrodes, amplified, and sent to the projector. Daniel Dennett describes an unexpected result of the experiment in his *Consciousness Explained*:

One might suppose that the patients would notice nothing out of the ordinary, but in fact they were startled by the effect, because it seemed to them as if the slide projector was anticipating their decisions. They reported that just as they were "about to" push the button, but before they had actually decided to do so, the projector would advance the slide—and they would find themselves pressing the button with the worry that it was going to advance the slide twice!

That odd effect, caused by the delay between the *decision* to do something and the *awareness* of that decision, raises profound questions about the nature of consciousness. But for the moment, let's just note that patients were able to control a useful machine with their brains alone, even if they didn't realize that's what they were doing.

Researchers are divided on the question of which is the better method for getting signals from the brain, implanted electrodes or EEG. Both techniques have adherents. Both also have shortcomings. Implants can detect the focused and precise electrical activity of a very small number of neurons, while EEG can only pick up signals *en masse* and distorted by the skull. EEG is noninvasive, while implanted electrodes require risky brain surgery. The two schools of thought coexist and compete peaceably for headlines and limited grant money, although there is some ill will between them and badmouthing occasionally surfaces in the press.

The EEG-based approach dates back at least to the late 1980s, when Emanuel Donchin and Lawrence Farwell, the erstwhile collaborators

now on opposite sides of the “brain-fingerprinting” controversy, devised a system that let test subjects spell with their minds. A computer would flash rows and columns of letters on a screen; when the row or column with the desired letter flashed repeatedly, a P300 wave was detected; this process was reiterated until the user had whittled the options down to one letter—and then the whole process would begin anew for the next letter. Donchin and Farwell found that their test subjects could communicate 2.3 characters per minute.

While that system clearly worked, it was indirect—that is, it relied on the uncontrollable P300 wave rather than on the user’s willful control of brain or machine. Most subsequent EEG-based researchers have sought more direct control. For example, the work of Gert Pfurtscheller, head of the Laboratory of Brain-Computer Interfaces at Austria’s Graz University of Technology, emphasizes the motor cortex, so the computer reacts when a subject imagines moving his extremities. A multinational European project, headed by Italy-based researcher José del Rocío Millán, has been working on a system called the Adaptive Brain Interface: the user’s brain is studied while he imagines performing a series of pre-selected activities (like picking up a ball); the brain pattern associated with each imagined activity then becomes a code for controlling a computer with one’s thoughts. Jonathan Rickel Wolpaw of the Wadsworth Center in the New York State Department of Health, the leading American authority on EEG-based brain-machine interfaces, told *Technology Research News* in 2005 that using his cursor-controlling system “becomes more like a normal motor skill”; the relationship between thought and action becomes even more direct.

The best-known European researcher who works on EEG-based brain-machine interfaces is University of Tübingen professor Niels Birbaumer. In 1995, he won the Leibniz Prize, a prestigious German award, for a successful neurofeedback therapy he devised to help epileptics control their seizures. With the prize money, he was able to fund his own research into the use of EEGs for brain-machine interfaces. He was soon testing what he called the “Thought Translation Device” on actual paralyzed patients—something many other researchers haven’t yet attempted with their brain-machine interfaces—and reported impressive successes not long after. One patient, a locked-in former lawyer named Hans-Peter Salzmann, was able, after months of training, to use the device to compose letters, including a thank-you note to Birbaumer published in *Nature* in 1999. In the following years, Salzmann’s system was connected to the Internet, so he could surf the Web and send e-mails. Here is how

Salzmann, in a 2003 interview with the *New Scientist* magazine, describes the mental gymnastics needed to control the cursor:

The process is divided into two phases. In the first phase, when the cursor cannot be moved, I try to build up tension with the help of certain images, like a bow being drawn or traffic lights changing from red to yellow. In the second phase, when the cursor can be moved, I try to use the tension built up in the first phase and kind of make it explode by imagining the arrow shooting from the bow or the traffic lights changing to green. When both phases are intensely represented in my head, the letter is chosen. When I want to not choose a letter, I try to empty my thoughts.

Although Birbaumer has reportedly had good results with some of the more than a dozen other patients he has worked with, none has been as successful as Salzmann, and even he has off-days.

Birbaumer's most astonishing case has been that of Elias Musiris, the owner of factories and a casino in Lima, Peru. ALS left Musiris totally locked in by the end of 2001, unable even to blink or control his eyes. A profile of Birbaumer in *The New Yorker* describes the scientist's visit with Musiris in the summer of 2002 and how, after several days of practice and training, Musiris was able to answer yes-or-no questions and to spell his own name with the Thought Translation Device. He had been unable to communicate for half a year. No fully locked-in patient—incapable even of blinking or eye motion—had ever communicated anything before.

Mind Over Matter

Birbaumer thinks the implant approach to brain-machine interfaces is less practicable than the EEG approach, even though the latter is slower. He says that his patients prefer sluggish communications over having a hole in the head.

But a few patients have said yes to a hole in the head, in hopes of controlling machines with their brains. The first were patients of Philip R. Kennedy, an Emory University researcher who, in the 1980s, invented and patented an ingenious new neural electrode. Even setting aside the many health risks of having an electrode surgically implanted in your brain, there were a host of technical problems associated with previous brain implants. Sometimes scar tissue formed around them, reducing the quality of the electrical signals they picked up. Sometimes the electrodes would shift within the brain, so they no longer picked up signals from the same neurons. Kennedy's new design solved some of these problems. The

tip of his electrode was protected in a tiny glass cone; once it is implanted, neurons in the brain actually grow into the cone and reach the electrode. The electrode is thus sheltered from scarring and jostling.

After experiments with rats and monkeys, Kennedy obtained FDA permission in 1996 to test his implant in human patients. The first patient, a woman paralyzed by ALS and known only by the initials M.H., could change the signals the electrode detected by switching her mental gears; there was a distinct difference between when she concentrated furiously and when she let her mind idle. Unfortunately, she died two and a half months after the surgery.

Kennedy's second patient was Johnny Ray, a Vietnam vet and former drywall contractor locked in by a stroke. He received his implant in March 1998, and over the next few months learned to move a cursor around a screen by imagining he was moving his hand. By the time the press was informed in October 1998, Ray was able to move the cursor across a screen with icons representing messages—allowing him to indicate hunger or thirst, and to pick from among messages like “See you later.” After months of further practice he was able to spell, using the cursor to hover toward his desired letter and then twitching his shoulder—one of the few residual muscles he could control—to select it, like clicking a computer mouse.

When asked what he felt as he moved the cursor, Ray spelled out “NOTHING.” This couldn't have been strictly true: it is clear that moving the cursor was exhausting work. But the doctors interpreted this to mean that Ray no longer had to imagine moving his hand. That intermediate step became unnecessary; he now just *thought* of moving the cursor and it responded.

Ray died in 2002, but Kennedy and his colleagues have carried on their work with several other patients. In his more recent studies, Kennedy has reportedly increased the number of electrodes he implants, giving him access to a richer set of brain signals. But the number of electrodes Kennedy implants is dwarfed by the number of electrodes on the implants used by the only other brain-machine interface researchers to put long-term electrodes into humans. That team, led by Brown University professor John P. Donoghue, recently obtained permission to conduct two clinical implant studies—one on paralyzed patients, the other on patients with motor neuron diseases like ALS. Their system, called BrainGate, uses 96 tiny electrodes arrayed on an implant the size of an M&M. Seen magnified, the implant looks like a bed of nails.

As of this writing, two patients have had BrainGate implants inserted in their heads. While only preliminary details have been released about

the second patient, the first patient's story has been widely publicized. Matthew Nagle was stabbed in the neck with a hunting knife during an altercation at an Independence Day fireworks show in 2001. His spinal cord was severed, leaving him quadriplegic. Although communication isn't a problem for him—he can talk, and he has given interviews and testified at his attacker's trial—he agreed to participate in the BrainGate study, and was surgically implanted in June 2004. The 96 electrodes in his head are estimated to be in contact with between 50 and 150 neurons, and signals from about a dozen have been used to give him the same sort of cursor control Johnny Ray had. Nagle's computer was also hooked up to other devices, so he could use it to change the volume on a television and turn lights on and off.

Several researchers have also done impressive work with electrodes in animals. The leaders in this field are unquestionably Duke University neurobiologist Miguel A.L. Nicolelis and State University of New York neurobiologist John K. Chapin. In 1999, they demonstrated that rats could control a robotic lever just by thinking about it. The rats had been trained to press a bar when they got thirsty; the bar activated a little robotic lever that brought them water. Electrodes in the rats' heads measured the activity of a few neurons, and the researchers found patterns that occurred whenever the rats were about to press the bar. The researchers then disconnected the bar, turning it into a dummy, and set up the robotic lever to respond whenever the right brain signals were present—much as Grey Walter had used a dummy remote control with his slide projector. Some of the rats soon discovered that they didn't have to press the bar, and they began to command the robotic lever mentally.

Nicolelis, Chapin, and their colleagues quickly extended the experiment, and within a couple of years reported successes in getting monkeys to control a multi-jointed robotic arm. To be precise, the monkeys didn't *know* they were controlling a robotic arm: they were trained, with juice as a reward, to use a joystick to respond to a sort of video game, while unbeknownst to them the joystick was controlling the robotic arm. Their brains' electrical signals were measured and processed and interpreted. The researchers then used the brain signals to control the robotic arm directly, turning the joystick into a dummy. Eventually the joystick was eliminated altogether. As a bit of a stunt, the researchers even sent the signals over the Internet, so that a monkey mentally controlled a robotic arm hundreds of miles away—unwittingly, of course.

Other scientists have improved and varied these experiments further still. Andrew B. Schwartz, a University of Pittsburgh neurobiologist who

has for more than two decades studied the electrical activity of the brains of monkeys in motion, has trained a monkey to feed itself by controlling a robotic arm with its mind. In video of this feat available on Schwartz's website, the monkey's own limbs are restrained out of sight, but a robotic arm, with tubes and wires and gears partially covered by fake plastic skin, sits beside it. A gloved researcher holds a chunk of food about a foot away from the monkey's mouth, and the arm springs to life. The shoulder rotates, the elbow bends, and the claw-hand takes the chunk of food, then brings it back to be chomped by the monkey's mouth. The researcher holds the chunk closer, and the monkey changes his aim and gets it again. The whole time, the back of the monkey's head, where the electronic apparatus protrudes, is discreetly hidden from view.

No one can deny that these are all breathtaking technical achievements. Neither should anyone deny that there are a number of major interlocking obstacles that must be overcome before implant-based brain-machine interfaces will be feasible therapeutic tools for the thousands of people who could, in theory, benefit from their use.

The first problem relates to implant technology itself. Implant design is rapidly evolving. Newer implants will have more electrodes; implants with thousands of electrodes will be tested in the next few years. New materials and manufacturing processes will allow them to shrink in size. And implants will likely become wireless. These advances will carry with them new problems to be solved; wireless implants, for example, might cause thermal effects that weren't a problem before.

Second, even though biocompatibility is always considered when designing and building brain implants, most implants don't work very well after a few months in a real brain. There are exceptions—electrodes in a few test animals have successfully picked up readings for more than five years—but in general, implant longevity is a problem. One way around it might be to use electrodes capable of moving small distances to get better signals (a notion proposed by Caltech researcher Richard A. Andersen in 2004).

Third, there is still much disagreement about which spots in the brain give the most useful signals for brain-machine interfaces. And much work needs to be done to improve the "decoding" of those signals—the signal processing that seeks to discern meaning in the measurements.

Finally, as the technology moves slowly toward commercial viability, standard practices, procedures, and protocols will have to be established, and there will be challenges from government regulators on issues like safety and consent.

In time, the technology will improve, and implant-based brain-machine interfaces will be worthwhile for a great many patients. But as things stand today, they make sense for almost no one. They involve significant risk. They are expensive, thanks to the surgery, equipment, and manpower required. They can be exhausting to use, they generally require a lot of training, and they aren't very accurate. Only a locked-in patient would benefit sufficiently, and even in some locked-in cases it wouldn't make sense. For all the *technical* research that has been done, there has been very little *psychological* research, and we still know very little about the wishes and aspirations of severely paralyzed patients.

Artificial Limbs

Experiments allowing animals to mentally move robotic arms raise the question: To what extent will brain-machine interfaces allow paralyzed humans to regain mobility?

One sure bet is that some paralyzed patients will be able to control their own hands and arms, at least in a rudimentary fashion. A little-known but remarkable technology that has been used clinically for more than two decades can restore very basic control to paralyzed muscles. The technology is called Functional Electrical Stimulation (FES). It uses electrical impulses, either applied to nerves or directly to muscle, to jumpstart paralyzed muscles into action. FES has become an important physical therapy tool for some paralytics, allowing them to exercise muscles they can't control. But it can do much more: It has been used to give paralyzed patients new control over their bladder and bowels; it has been used to help several hundred paraplegics stand and haltingly walk with a walker; it has even been used in a number of cases to give quadriplegics a semblance of control over their arms and hands. The first patient to use FES to control his own hands was Jim Jatich, a design engineer left quadriplegic by a 1978 diving accident. In 1986, he had stimulating electrodes implanted into his hands; he can control those implants with a sort of joystick technology manipulated by his chin. Thanks to this system, Jatich and hundreds like him can use computers, write with pens, groom themselves, and eat and drink on their own.

It takes no great leap of the imagination to see how this approach might work in conjunction with the cursor-controlling systems, and indeed, researchers at Case Western Reserve University reported in 1999 that they had already combined the two technologies. A test subject who used FES to open and close his disabled hand was first trained to move a

cursor using an EEG-based brain-machine interface. Then the EEG signal was connected to the FES, so that when he controlled his brain waves he could open and close his hand. A more recent study by researcher Gert Pfurtscheller used a similar approach, finding that a patient who triggered his FES by changing his EEG activity “was able to grasp a cylinder with the paralyzed hand.” Researchers using brain implants have taken notice, too. “Imagine if we could hook up the sensor directly to this FES system,” implant pioneer John Donoghue told *The Scientist* in 2005. “By thought alone these people could be controlling their arm muscles.”

But FES doesn’t work for everyone. Patients with many kinds of nerve and muscle problems can’t use FES—and, needless to say, amputees cannot use it either. Such patients might instead turn to robotics. Donoghue has already shown that his BrainGate system can be used for basic robotic control: His patient Matthew Nagle was able to open and close a simple robotic hand using his implant. That sort of robotic hand is increasingly available to amputees, replacing the older mechanical prostheses normally controlled by cables. The newer robotic prostheses are usually controlled by switches, or by the flexing and flicking of muscles in the amputee’s stump. And some more advanced models respond to electromyographic activity—that is, the electrical activity in muscles.

Consider the case of Jesse Sullivan. A bespectacled average Joe in his fifties, Sullivan was fixing electrical lines for a Tennessee power company in 2001 when he was badly electrocuted. Both his arms had to be amputated and he was fitted with mechanical prostheses. Then, researchers led by Todd A. Kuiken of the Rehabilitation Institute of Chicago replaced Sullivan’s left prosthetic arm with a robotic arm he can control through nerves grafted from his shoulder to his chest. This lets him move his robotic arm just by thinking where he wants it to go, and according to the institute’s website, “today he is able to do many of the routine tasks he took for granted before his accident, including putting on socks, shaving, eating dinner, taking out the garbage, carrying groceries, and vacuuming.”

It will be many years before any locked-in patient can control a robotic limb that fluidly. The brain-machine interfaces that let patients slowly and sloppily move a cursor today might be able to control a simple and clunky claw, but nothing that matches the complexity of Jesse Sullivan’s new arms. And even Sullivan’s high-tech robotic limbs don’t come near to rivaling the versatility of the real thing. A real human arm has seven degrees of freedom and a hand has twenty-two degrees of freedom. While robotic limbs will surely be built with that level of complexity, capable of imitating (or surpassing) all the billions of positions that a human arm

and hand can take, it is hard to see how such complex machines can ever be controlled by either the muscle signals that Jesse Sullivan uses or by a descendant of today's brain-machine interfaces. There is just too much information required for dexterous control. Born as infants "wired," so to speak, with countless neuronal connections in our limbs, it takes us years to master our own bodies. No artificial appendage will get that intimate and intricate a connection.

Of course, paralyzed patients and amputees don't necessarily need full equivalency; even partial functionality can dramatically improve their quality of life. And there is no reason why a patient would have to control every aspect of a prosthetic limb—some of the mental heavy-lifting could be done by computers built into the prosthesis itself. So while a patient might use a brain-machine interface to tell an artificial hand to grasp a cup, the hand itself might use computerized sensors to tweak the movements and adjust the firmness of the grasp. As one of the researchers already working on this concept told *The Scientist*, this "shared control" idea "seems to make the tasks a lot more reliable than having solely brain-based control."

This concept can be extended even further. Wheelchairs controlled by EEG or brain implants are plausible—although patients using breath-control devices can operate their wheelchairs more adroitly than would be possible with any of today's brain-machine interfaces. And if brain-controlled wheelchairs might someday be available, why not other machines? Several research teams around the world have been working for years on exoskeletons. While these robotic suits are generally intended for use by soldiers, the elderly, or the disabled, there are many other possible applications, as a recent article in *IEEE Spectrum* points out: "Rescue and emergency personnel could use them to reach over debris-strewn or rugged terrain that no wheeled vehicle could negotiate; firefighters could carry heavy gear into burning buildings and injured people out of them; and furniture movers, construction workers, and warehouse attendants could lift and carry heavier objects safely." Making exoskeletons work with brain-machine interfaces for severely paralyzed patients is a distinct, if distant, possibility.

The Higher Senses

Unsurprisingly, much of the funding for research on brain-machine interfaces has come from the United States military, to the tune of tens of millions of dollars. Most of this funding has come through DARPA, the Pentagon's bleeding-edge R&D shop, although the Air Force and the

Office of Naval Research have also chipped in substantially. (DARPA's British and Canadian equivalents have also, to a lesser extent, funded brain-machine interface work over the years.) DARPA's interest in robotics and brain-machine interfaces is quite broad—according to its website, the agency would like to find ways to “seamlessly integrate and control mechanical devices and sensors within a biological environment.”

DARPA is also interested in a less sophisticated form of mental control over military aircraft—one in which aircraft are made more responsive to the needs and wishes of pilots and aviators by closely monitoring them with sensors and adapting accordingly. This intriguing approach—given names like the “cognitive cockpit” and “augmented cognition” (augcog)—would rely on EEG and other indicators. An *Aviation Today* article explains it this way: “Instead of merely reacting to pilot, sensor, and other avionics inputs, the avionics of tomorrow could detect the pilot's internal state and automatically decrease distractions, declutter screens, cue memory or communicate through a different sensory channel—his ears vs. his eyes, for example. The system would use the behavioral, psychophysiological, and neurophysiological data it collects from the pilot to adapt or augment the interface to improve the human's performance.”

It will be years before that sort of adaptive cockpit is regularly implemented. And even that is a far cry from the idea of direct mental control of airplanes. This is sometimes called the “Firefox” scenario, after the mediocre 1982 action flick *Firefox*, in which Clint Eastwood was ordered to steal a shiny new Soviet fighter jet specially rigged to read and obey the pilot's mind so as to save him the milliseconds it would take to actually press buttons. There's a hilarious catch, though: the brain-reading technology, which is built into Eastwood's flight helmet, can only read thoughts mentally expressed *in Russian*. At the film's climax, Eastwood must destroy a pursuing fighter, but the mission is almost ruined when he forgets to mentally fire his missiles in Russian. At the last moment, he remembers, the missiles fire, and the day is saved. In real life, it is hard to see why brain-piloting of a fighter jet would ever be necessary or desirable, especially given the advances in unmanned aircraft controlled by computers or by remote humans.

If brain-machine interfaces are to advance sufficiently for people to control robots with their minds or for the severely disabled to interact normally with the world around them, researchers will have to improve not just the ability to detect and decode the brain's commands, they will also have to improve the feedback that users get. A locked-in patient moving a cursor on a screen can see the results of his mental exertions, but it

would be much harder for him to tell, for example, how tightly a robotic hand is grasping a Fabergé egg.

Normally, our senses give our brains plenty of feedback, especially our senses of hearing, vision, touch, and proprioception (balance and orientation). For patients who are deaf, blind, or disabled, researchers and therapists have long sought methods by which one sense could be substituted for another. Haptic technology, by which sensory information is translated into pressure on the skin, has been around for decades; it is central to telerobotics and it has even been used to give some blind patients a very crude kind of “vision” by translating camera images into tactile sensations. It is also of consummate interest to researchers and theorists working on virtual reality. Some basic version of haptic technology, one that puts pressure somewhere on a locked-in patient’s skin, would be a simple way to give at least a little non-visual feedback for controlling robotic devices.

And what if it were possible to create the illusion of tactile sensation by directly stimulating the nervous system? Such illusions have long been created haphazardly by researchers stimulating nerves and the brain with electrodes, but what if they could be produced in an organized way, to correspond with the motions of a prosthetic device? A team led by University of Utah bioengineering professor Kenneth W. Horch has reportedly been able to do just that. They ran wires from a robotic arm to nerves in the forearms of amputees. The wires sent nerve signals to the robotic arm, giving the amputees control over the robot. But the wires also carried electrical impulses back into the amputees’ nerves, giving them feedback from the arm. According to a 2004 article in *The Economist*, this enabled the patients “to ‘feel’ natural sensations as though through the device’s fingers,” and “made it possible for them to gauge how much pressure to apply when commanding the motors to grip. In addition, position sensors in the robot’s joints were translated into ‘proprioception’ signals that enabled the subjects to feel the arm’s position, even when their eyes were closed.” This technology is still quite far from practical use, but it is a very impressive proof-of-concept.

There have also been advances in the electric creation of other perceptions. Cochlear implants, devices which can restore hearing in patients with certain kinds of deafness or hearing loss, have constantly improved since they went on the market in the early 1980s. These implants, which circumvent the natural hearing mechanism by electrically stimulating nerves in the ear, have now been used by more than 80,000 patients. In those rarer cases of deafness caused by damage to the auditory nerve,

cochlear implants are not an option; in some of these cases—more than 300 around the world so far—researchers have begun putting implants directly on the auditory brainstem. The sound quality of auditory brainstem implants is greatly inferior to even that of cochlear implants, but still valuable to those with no other options.

To a lesser extent, there has been progress in artificial vision. Several research teams have been working on implants that could replace *some* of the functionality of *some* of the retina of patients whose own retinas are degenerating. Each team has its own approach. Some researchers are working with simple implants that go just behind the retina, using sunlight to stimulate degenerating retinal cells; they need no wires or batteries. Others have been experimenting with bulkier systems that send data from external cameras to chips implanted in front of the retina. In clinical trials in the past few years, both approaches have succeeded in improving vision in a handful of patients—very minimal improvements, but improvements nonetheless. And for patients whose vision problems are unrelated to the retina, researchers have been working on tapping into the optic nerve or directly stimulating the visual cortex. One of the pioneers of the latter technique, maverick scientist William Dobelle, died in 2004; even though he had some successes with patients, it isn't clear whether any other researcher will emulate either Dobelle's technique (sending signals from cameras to the surface of the visual cortex) or his style (taking patients to Portugal so as to skirt the FDA approval process).

It must be emphasized that artificial vision research is today utterly primitive. Human vision is made possible by an almost unimaginably complex biological system involving millions of photoreceptors, billions of cells, and methods for processing information that researchers can still only guess at. Likewise, the technologies for controlling robotic limbs and restoring hearing will take many more years to mature. Our ignorance is awesome. But we have made progress; real patients have had lost powers restored to them. In our lives, the blind have received their sight, the lame have walked, and the deaf have heard. These are miraculous times.

The New Brain Science

The future of brain-machine interfaces will depend, in part, on several overlapping areas of research now in their infancy. Some of them will fade from view, forgotten footnotes even historians will ignore. Others will likely rise in importance in the years to come, shaping how we think about minds and computers.

First, the study of neurobiology and neurochemistry is progressing rapidly, and scientists are learning ever more about how the brain and nervous system function. The full import of some of the revolutions in brain science is only now beginning to be understood. Chief among these revolutions is the overthrow of the notion of the static brain. Thirty years ago, scientists believed the adult brain was “hardwired”—immutably fixed like electronic circuits. They now know that the brain is flexible, adaptive, and resilient. Understanding the extent of this *plasticity*—the ability of neurons to form new connections and to strengthen or weaken existing connections—is central to understanding how our brains grow, heal, and age.

A second area of research might make brain-machine interfaces unnecessary for some patients: biological solutions for some kinds of paralysis. Although brain damage due to strokes can’t be undone, and no magical stem-cell-derived cures for spinal cord injuries should be expected anytime soon, it is worth remembering that biomedical research is progressing simultaneously with the research on brain-machine interfaces. Indeed, the last few years have seen remarkable advances in growing, moving, and manipulating nerve cells in ways that could benefit paralyzed patients.

A third subject that interests scientists is a new way of manipulating the brain called Transcranial Magnetic Stimulation (TMS). Originally developed in the 1980s as part of the growing arsenal of brain-mapping techniques, the fundamental idea in TMS is this: Because the activity of the brain and nervous system is electrochemical in nature, very intense magnetic fields can alter the way the brain functions. This may sound like the sort of charlatanry that has been around for ages—from Mesmer’s *baquet* to the “magnetotherapy” cures advertised in late-night infomercials—but TMS is the real deal. Scientists have found that TMS can alter mood. It can reduce hallucinations and treat some migraines. It has been used by one researcher to generate an illusion of what he calls “sensed presence,” which he hypothesizes might explain the paranoid and the paranormal. Another researcher has shown that TMS can improve creativity, although only in some people and only very slightly and briefly; he also theorizes that by temporarily disabling some of our normal brain functions, TMS might be used to turn us into temporary savants. DARPA is apparently considering TMS as a tool to help soldiers perform well without sleep, and TMS is being clinically tested as a replacement for electroshock in the treatment of depression.

Investigators are also exploring other noninvasive techniques for mind-influence. In just the last two years, studies have shown that a specialized technique related to MRI called echo-planar magnetic resonance

spectroscopic imaging (EP-MRS) can temporarily improve the mood of patients with bipolar disorder. Another technique called direct current polarization—the use of a battery to send a very tiny electrical current through the front of the head—can slightly improve verbal ability, according to researchers.

As with electroshock therapy and Deep Brain Stimulation, no one is certain how TMS, EP-MRS, and direct current polarization produce the effects they do. It still isn't clear whether they can affect neurons deep in the brain or just those closer to the surface. Although none of these techniques is painful, little is known about their health risks or long-term effects. And these outside-the-skull techniques are blunt; they don't give fine control over the mind; they're more like a cudgel than a scalpel. But, taken together, these developments may presage a new interest in the use of machines to influence the mind.

A fourth area of research involves the study of disembodied brains and neurons in combination with silicon chips. The gory aspects of some of this research seem intended chiefly to attract the attention of editors and headline writers. One can be forgiven for wondering what real scientific value there is in removing the brains from a lamprey (an eel-like fish) and connecting them to a little robot. Or in wiring up a mass of disembodied rat brain cells to a robotic arm holding colored markers, so that the ex-brain blob's electrical activity is turned into "art." Or in the simple experiment that prompted absurd headlines like "Brain Grown from Rat Cells Learns to Fly Jet."

When you look behind the hype, some of this research has serious scientific value: it could improve our understanding of neurochemical processes, it could teach us about how neurons interact, and it could help in the design of better electrodes. It isn't obvious where this gruesome combining of silicon and neurons will lead, but it is clearly a growing area of research that cuts across several disciplines.

Fifth, and finally, a few scientists believe that computer chips could perform some of the higher functions of the brain. As we have seen, the vast majority of the research on brain-machine interfaces and neural prosthetics relates to the body's motor functions and sensory systems. But what if an implant could assist in the brain's cognitive functioning?

To date, the only serious effort to create a "cognitive prosthesis" is the work of a team of researchers headed by University of Southern California professor Theodore W. Berger. They are attempting to create a computer chip that can do some or all of the work of the hippocampus, part of the brain critical for the formation of long-term memories.

Damage to the hippocampus has been connected to amnesia; degeneration of the hippocampus is associated with Alzheimer's disease. As Berger and his colleagues describe their hopes in a recent article in *IEEE Engineering in Medicine and Biology*, if artificial "neurons" can approximate the functions of biological neurons, and ultimately replace damaged neurons, then we will see the rise of "a new generation of neural prostheses" that "would have a profound impact on the quality of life throughout society; it would offer a biomedical remedy for the cognitive and memory loss accompanying Alzheimer's disease, the speech and language deficits resulting from stroke, and the impaired ability to execute skilled movements following trauma to brain regions responsible for motor control."

These are grand ambitions. But Berger and his team have planned a gradual research program, starting by "reverse-engineering" the hippocampus—thoroughly analyzing the electrical functions of thin slices of rat brain—and then moving on to designing and testing microchips that can replicate those functions. Eventually those chips will be connected to living animal brains for testing.

Will Berger's approach work? And if so, will it someday lead to cognitive prostheses capable not only of restoring damaged brains but of doing much more, such as connecting brains telepathically or giving us editorial control over our memories? Speculation abounds, sometimes rooted in fantasies of the imagination rather than in the best scientific evidence.

Beyond the Cyborg

For the past many decades, serious brain-machine science has evolved alongside popular dreams and nightmares about the meaning of merging men and machines. These visions of the future make incremental advances in the laboratory seem like the slow march toward an inevitable age of cyborgs.

In the summer of 1947, the brilliant American mathematician Norbert Wiener coined the term *cybernetics*—derived from the Greek for "steersman"—to describe the study of "control and communication theory, whether in the machine or the animal." He considered cybernetics to be a vitally important new discipline, and he explained it in two books (*Cybernetics*, 1948, and *The Human Use of Human Beings*, 1950) that are surprisingly humanistic, especially in light of the subject matter and the author's impeccably technocratic credentials. For a short while, cybernetics aroused significant academic interest—at the intersection of physiology, computers, engineering, philosophy, economics, psychology,

and sociology. But eventually, its ideas were so fully absorbed into these disciplines that much of cybernetics came to seem obvious.

In 1960, at the height of interest in cybernetics, the word *cyborg*—short for “cybernetic organism”—was coined by researcher Manfred E. Clynes in a paper he co-wrote for the journal *Astronautics*. The paper was a theoretical consideration of various ways in which fragile human bodies could be technologically adapted and improved to better withstand the rigors of space exploration. (Clynes’s co-author said the word *cyborg* “sounds like a town in Denmark.”) Around the same time, Jack E. Steele, a polymath doctor-engineer-neuroanatomist serving in the U.S. Air Force, coined the word *bionics* for the use of principles derived from living systems to solve engineering and design problems.

These words and concepts soon entered the popular imagination, starting with a 1972 science fiction novel, *Cyborg*, that became the basis for the seventies TV show *The Six Million Dollar Man* about the “world’s first bionic man,” and then the spin-off *The Bionic Woman*. This was followed by decades of movies, TV shows, books, comics, and video games with cops, criminals, soldiers, and aliens who were cyborgs—from Darth Vader in the seventies to RoboCop in the eighties to the Borg in the nineties. The flood of cyborgs in pop-culture caught the attention of academics, and soon anthropologists, philosophers, and literary theorists were offering up unreadable piles of “cyborg scholarship.”

In popular usage, the term “bionic” now refers to any kind of electronic implant or prosthesis, and so several different people—including amputees with robotic arms, like Jesse Sullivan—have been dubbed “the world’s first” bionic man or woman. Similarly, the term “cyborg” has been overextended to the point of meaninglessness. Who was the first human cyborg? Maybe it was Johnny Ray, the first brain-machine interface implantee to control a computer cursor with his thoughts. Or maybe the Australian performance artist Stelios Arcadiou—called STELARC—known for his decades of grisly forays into high-tech body modification. Or perhaps Steve Mann, the Canadian wearable computer pioneer, who since the 1980s has spent most of his waking hours viewing the world through little screens in front of one or both of his eyes. Or Kevin Warwick, the professor in England whose audacious showmanship in having chips implanted in his body has brought him tremendous publicity despite the total lack of scientific merit in his stunts.

But to the most ambitious and most radical advocates of merging brains and machines, such advances are mere child’s play. These so-called transhumanists long for an age when human beings will leave the miser-

ies and limits of the body behind, and achieve new ecstasies of freedom and control: We will send feelings and conscious thoughts directly from brain to brain; we will delete unwanted memories at will; we will master complex subjects by “downloading” them directly into our minds; we will “jack in” to virtual realities; and eventually, we will be able to “upload” our personalities into computers or robots, where the self could live on indefinitely.

Such fantasies are staples of much science fiction, including cyberpunk novels like William Gibson’s *Neuromancer* and Neal Stephenson’s *Snow Crash*, and movies like *Total Recall*, *Johnny Mnemonic*, *Strange Days*, *The Final Cut*, and the Matrix trilogy. These books and movies are mostly dystopian visions of the future, or tales in which things have gone terribly awry—crime, cruelty, and mass delusion. But advocates of transhumanism, like Ramez Naam, author of the recent book *More Than Human*, are far more optimistic:

With neural prosthetics, information from the emotional centers of someone else—say, a loved one—could be piped straight to your empathy center. So rather than having to guess what your spouse or child is feeling, you would simply be sensing it via the wireless link between your brains.... The end result might be just like having an unusually keen sense of how others are feeling, with the option to dial that sense up or down in intensity according to whatever criteria you choose.

Naam imagines how that sharing of feelings might be taken further:

You send your spouse what you see and hear and feel.... That night, you and your spouse make love while opening all of your senses and emotions to each other. The intimacy is beyond anything you have known.

And further still:

In principle we could do this for all the senses—record not just what you see, but also what you hear, taste, smell, and feel, all at the level of your brain. Playing back such an experience would be a little like reliving it. You might even be able to play that kind of sensory recording back for someone else, turning the experience you had into a set of nerve impulses that could be sent into the other person’s brain, allowing him or her to experience at least the sensory parts of an event from your perspective....

When sensations, emotions, and ideas become digital, it’s as easy to share them with a dozen friends, or a thousand strangers, as it is to

send them to one person.... We'll be able to broadcast the inner states of our minds.

What an unattractive vision of the future—this world in which you can snoop on your children's feelings, feel what it's like to have sex with yourself, and broadcast the full sensory experience of your sexual encounters to the world. These are shallow, solipsistic aspirations, utterly divorced from the hopes and fears of mature human beings. The transhumanist fantasy is surely not the best guide for thinking about the genuine ethical dilemmas we now face at the dawn of the Age of Neuroelectronics.

A True Humanism

Without question, there are genuine human benefits to be gained if brain-machine technology advances in a sober, limited way. People with motor diseases or severe mental disorders can be helped with brain implants, amputees might find a new freedom and mobility in the use of mind-controlled prosthetics, the blind might get new electronic eyes and the deaf new ears, and even severely paralyzed patients might someday be "unlocked." Yet it is also possible to envisage a world where these new technologies are used for less noble purposes—from the next generation flight into alternative reality to the active manipulation and control of innocent subjects to the self-destructive pursuit of neurological perfection.

In the short term, brain implants probably shouldn't worry us: they are still very crude and the risks of brain surgery make them worthwhile only for those without other options. And we probably should not exert much energy fretting about the transhumanist future, which requires a level of scientific sophistication so far removed from the present that making predictions about its plausibility is a fool's errand. The greatest questions lie in the middle range—that time, some years hence, when today's techniques are vastly improved, when brain surgery becomes safe enough and implants become effective enough for the electronic alteration of the brain to evolve from desperate therapy to mainstream enhancement. Not long ago, the prospect of manipulating our minds with machines would have been universally disquieting. Now, after decades of "softening up" by advances in science and science fiction, far fewer people find the notion of neuro-enhancement troublesome. Its potential clients are not just the radicals who long for a post-human future, but ordinary people who grew up in an age of transplants and implants, of fictional bionic men and vivid cyborg fantasies.

The obvious temptation will be to see advances in neuroelectronics as final evidence that man is just a complex machine after all, that the brain

is just a computer, that our thoughts and identity are just software. But in reality, our new powers should lead us to a different conclusion: even though we can make the brain compatible with machines to serve specific functions, the thinking being is a being of a very different sort. We are neither machines nor ghosts, but psychophysical unities—finite yet creative, embodied yet spiritual, cognitive yet not cognitive alone. No machine, however sophisticated, seems likely to duplicate or surpass that improbable mix of excellence, depravity, dignity, and uncertainty that is our human lot. On this score, the machine makers of the future still have much to learn from the myth makers of the past. And even as we seek to improve human life by improving the brain, we should be careful not to make ourselves into something worse in the effort to become something better.